arXiv:1703.09513v1 [cs.AI] 28 Mar 2017

Mining Best Closed Itemsets for Projection-antimonotonic Constraints in Polynomial Time

Aleksey Buzmakov[1] Sergei O. Kuznetsov[1]
Amedeo Napoli[2]
[1]National Research University Higher School of Economics
Moscow / Perm, Russia
[2] LORIA (CNRS – Inria – University of Lorraine)
Vandœuvre-lès-Nancy, France
avbuzmakov@hse.ru, skuznetsov@hse.ru, amedeo.napoli@loria.fr

March 29, 2017

Abstract

The exponential explosion of the set of patterns is one of the main challenges in pattern mining. This chalenge is approached by introducing a constraint for pattern selection. One of the first constraints proposed in pattern mining is support (frequency) of a pattern in a dataset. Frequency is an anti-monotonic function, i.e., given an infrequent pattern, all its superpatterns are not frequent. However, many other constraints for pattern selection are neither monotonic nor anti-monotonic, which makes it difficult to generate patterns satisfying these constraints. In order to deal with nonmonotonic constraints we introduce the notion of "projection antimonotonicity" and $\Sigma o \varphi \iota \alpha$ algorithm that allow generating best patterns for a class of nonmonotonic constraints. Cosine interest, robustness, stability of closed itemsets, and the associated Δ-measure are among these constraints. $\Sigma o \varphi \iota \alpha$ starts from light descriptions of transactions in dataset (a small set of items in the case of itemset description) and then iteratively adds more information to these descriptions (more items with indication of tidsets they describe). In the experiments, we compute best itemsets w.r.t. some measures and show the advantage of our approach over postpruning approaches.

1 Introduction

Interestingness measures were proposed to overcome the problem of combinatorial explosion of the number of valid patterns that can be discovered in a

dataset [VT14]. For example, pattern support, i.e., the number of transactions covered by the pattern, is one of the most famous measures of pattern quality. In particular, support satisfies the property of anti-monotonicity (aka "a priori principle"), i.e., the larger the pattern is the smaller the support is [MTV94, AS94]. Many other measures can be mentioned such as pattern stability [Kuz07, ROK08], margin closeness [MTU11], MCCS [SDB13], cosine interest [CWW14], pattern robustness [TMC14], etc.

Some of these measures (e.g., support, robustness for generators [TMC14], or upper bound constraint of MCCS [SDB13]) are "globally anti-monotonic", i.e., for any two patterns $X \sqsubseteq Y$ (where $\sqsubseteq$ stays for containment or subsumption relation in the pattern language) we have $\mathcal{M}(X) \geq \mathcal{M}(Y)$, where $\mathcal{M}$ is a measure. When a measure is anti-monotonic, it is relatively easy to find patterns whose measure is higher than a certain threshold (e.g. patterns with a support higher than a threshold). In contrast some other measures are called "locally anti-monotonic", i.e., for any pattern X there is an immediate subpattern $Y \prec X$ such that $\mathcal{M}(Y) \geq \mathcal{M}(X)$. The corresponding constraint induces an accessible system [BHPW10] in itemset data. Indeed, for any itemset selected by a locally anti-monotonic constraint, one can always find a smaller selected itemset different only in one item. The good strategy in this case is extension of a pattern Y only to patterns X such that $\mathcal{M}(X) \leq \mathcal{M}(Y)$. For example, cosine interest [CWW14] is "locally anti-monotonic," some other examples can be found in [BHPW10].

The most difficult case is when a nonmonotonic measure is not even locally anti-monotonic. The valid patterns can be selected by postpruning, i.e., by finding a (large) set of patterns satisfying an anti-monotonic constraint and pruning them w.r.t. the chosen nonmonotonic measure [ROK08, MTU11, TMC14]. For that one can rely on certain heuristics such as the one used in leap search [YCHY08]. More elaborated approaches allow constructing a measure from (anti-)monotonic primitives [SC05, CRB08]. These approaches find a good anti-monotonic relaxation of the measure for the dataset in hand. Another interesting approach for dealing with non-monotonic constraints is search for a closure operator on the set of patterns adequate for the constraint in question [SC08].

In this paper we deal with a recently introduced algorithm Σοφια, i.e. Sofia, for "Searching for Optimal Formal Intents Algorithm". Σοφια was applied for an interval-tuple data [BKN15]. In this paper we apply Σοφια for extracting the best itemsets w.r.t. a wide class of constraints. We introduce the polynomial version of the algorithm by accordingly adjusting the threshold and deeply studying the properties of the involved measures. Our algorithm is applicable to a class of measures called "projection-antimonotonic measures" or more precisely "measures anti-monotonic w.r.t. a chain of projections". This class includes globally anti-monotonic measures such as support, locally anti-monotonic measures such as cosine interest and some of the nonmonotonic measures such as stability or robustness of closed patterns. We should notice that this class of measures is not covered by the previously introduced approaches. In particular, for the primitive-based approaches [SC05, CRB08] it is not clear how one can

express certain measures from our class, e.g., stability and robustness, by means of the primitives. On the other hand, the approach for finding adequate closure [SC08] could be applied for stability and robustness, but the number of classes of equivalences that should be enumerated is likely to be high and accordingly the efficiency of the approach is likely to be low. Furthermore, neither of these approaches ensure a polynomial complexity of the algorithm.

In the experimental part of the paper we show that Σοφια can be efficiently used to mine itemesets w.r.t. a constraint based on Δ-measure, a polynomially computable analog of stability and robustness. It significantly outperforms the postpruning approaches based on best known algorithms for mining closed itemsets. We should mention that comparison of Σοφια with primitive-based approaches or with the approach for finding adequate closure is not possible since it requires a heavy study of efficiently expressing stability and robustness in terms of the primitives.

In the rest of the paper we work with itemsets and accordingly we use the word 'itemset' instead of 'pattern'. The remainder of the paper is organized as follows. Since the lattice of closed itemsets (concept lattice) is of high importance for concise representation of itemsets [PBTL99], we use the language of Formal Concept Analysis (FCA) [GW99] and pattern structures [GK01] which are introduced in Section 2. Then, Σοφια algorithm is detailed in Section 3 for projection-antimonotonic measures. In the next section we discuss cosine interest, robustness, and stability that are examples of such measures. Experiments and a discussion on Σοφια efficiency are proposed in Section 5, before the conclusion.

2 Preliminaries

2.1 Binary Dataset

FCA is a very convenient formalism for describing models of itemset mining and knowledge discovery [GW99]. Since [PBTL99] lattices of closed itemsets (concept lattices) and closed descriptions are used for concise representation of association rules. FCA gives a formalism for itemset mining. For more complex data such as sequences, graphs, interval tuples, and logical formulas one can use an extension of the basic model, called pattern structures [GK01]. With pattern structures one defines closed descriptions that give a concise representation of association rules for different types of descriptions with a partial order of "part-whole" (e.g., subgraph isomorphism order) or "is a" (e.g., "class-subclass") giving rise to a semilattical similarity operation $\sqcap$ [KS05, KKN11].

A *binary dataset* is a triple $\mathbb{D} = (T, \mathcal{I}, R)$, where T is a set of transaction identifiers, $\mathcal{I}$ is a set of items and $R \subseteq T \times \mathcal{I}$ is incidence relation giving information about items related to every transaction. A pattern structure or a (general) dataset is a triple $(T, (D, \sqcap), \delta)$, where $(D, \sqcap)$ is a semilattice of "descriptions" with similarity operation $\sqcap$ inducing natural partial order $(D, \sqsubseteq)$ given by $x \sqsubseteq y \Leftrightarrow x \sqcap y = x$ and $\delta : T \to D$ is a mapping from transactions

	i_1	i_2	i_3	i_4	i_5	i_6
t_1	x		x			
t_2		x	x			
t_3			x	x		
t_4			x		x	
t_5						x

(a) A binary dataset.

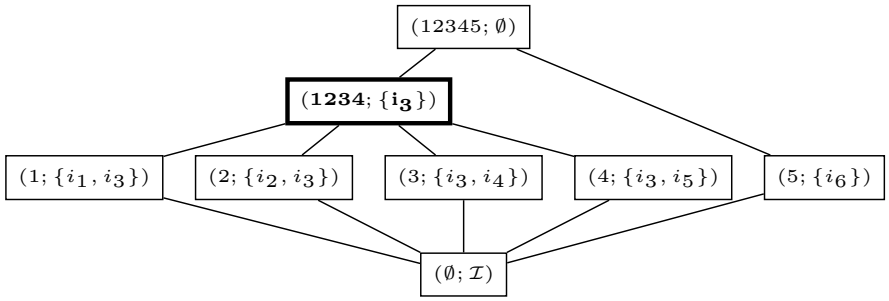

(b) A concept lattice. Concept extents are given by their indices, i.e., $\{t_1, t_2\}$ is given by 12.

Figure 1: A binary dataset and the corresponding concept lattice.

to their descriptions. Then, a binary dataset is $(T, (2^{\mathcal{I}}, \cap), \delta)$, where $\sqcap$ is $\cap$, $\sqsubseteq$ is $\subseteq$, and $\delta(t) = \{i \in \mathcal{I} \mid (t, i) \in R\}$. We use the pattern structure representation in order to iteratively modify the pattern space which is discussed later. Any subset of T is called a tidset and any subset of $\mathcal{I}$ is called an *itemset*. An example of a dataset is given in Figure 1a.

The following mappings give a Galois connection between the powerset of transactions and the semilattice of descriptions, e.g. $(2^{\mathcal{I}}, \cap)$.

$$d(A) := \bigsqcap_{t \in A} \delta(t), \qquad \text{for } A \subseteq T$$

$$t(x) := \{t \in T \mid x \sqsubseteq \delta(t)\}, \qquad \text{for } x \in D$$

In case of a binary dataset, the mapping $d(A)$ returns the maximal itemset common to all transactions in A, while the mapping $t(x)$ returns the set of all transactions whose descriptions are supersets of x. One can define closure operators and the corresponding closed tidsets and closed descriptions: $c_t := t \circ d$ and $c_d := d \circ t$ are closure operators, while the closed tidset A and closed itemset x are given by $c_t(A) = A$ and $c_d(x) = x$. As stated in [PBTL99], this type of closure (based on Galois connection and $\sqcap$ operation) is equivalent for itemsets to the closure wrt. "counting inference", which is common in data mining. However, the former definition unifies very important notions of "maximal common part", closure, and lattices of closed patterns, so we shall keep to it in this paper.

A *concept* of a dataset $(T, (D, \sqcap), \delta)$ is a pair (A, x), where $A \subseteq T$, called *extent* and $x \subset \mathcal{I}$, called *intent*, such that $d(A) = x$ and $t(x) = A$. In this case both A and x are closed tidset and itemset, respectively. The set of concepts is partially ordered w.r.t. inclusion on extents, i.e., $(A_1, x_1) \leq (A_2, x_2)$ iff $A_1 \subseteq A_2$ (or, equivalently, $x_2 \sqsubseteq x_1$), forming a lattice. An example of a lattice corresponding to the binary dataset in Figure 1a is given in Figure 1b.

In the reminder we need some results from pattern structures for justifying our approach. Moreover, our approach is also applicable to more complex data given by general datasets (pattern structures). For example, Σοφια was successfully applied to interval-tuple datasets [BKN15].

2.2 Projections of Datasets

The approach proposed in this paper is based on projections introduced for reducing complexity of computing with pattern structures [GK01].

A *projection* $\psi : D \to D$ is an "interior operator", i.e., it is (1) monotone ($x \sqsubseteq y \Rightarrow \psi(x) \sqsubseteq \psi(y)$), (2) contractive ($\psi(x) \sqsubseteq x$) and (3) idempotent ($\psi(\psi(x)) = \psi(x)$). A *projected dataset* $\psi(\mathbb{D}) = \psi((T, (D, \sqcap), \delta))$ is a dataset $(T, (D_\psi, \sqcap_\psi), \psi \circ \delta)$, where $\psi(D) = \{x \in D \mid \exists x^* \in D : \psi(x^*) = x\}$ is the fixed set of ψ and $\forall x, y \in D, x \sqcap_\psi y := \psi(x \sqcap y)$.

In the case of binary datasets projections correspond to removal of some items, with the respective change of the dataset $(T, (2^{\mathcal{I}}, \cap), \delta)$. The projection of an itemset $X \subseteq \mathcal{I}$ corresponding to removal of a set of items $Y \subseteq \mathcal{I}$ is given by

$$\psi(X) = X \cap (\mathcal{I} \setminus Y) = X \setminus Y. \tag{1}$$

Given a projection ψ we call $\psi(D) = \{x \in D \mid \psi(x) = x\}$ the *fixed set of* ψ. The fixed set contains those itemsets that contain no items from the set Y (the set of removed items). The projections are ordered w.r.t. inclusion of the fixed points (or by inclusion of the sets of removed items in the case of binary data), i.e., $\psi_1 < \psi_2$, if $\psi_1(D) \subseteq \psi_2(D)$, we say that ψ_1 is simpler than ψ_2 or that ψ_2 is more detailed than ψ_1.

Our algorithm is based on this order on projections. The simpler a projection ψ is, the less itemsets we can find in $\psi(\mathbb{D})$, and the less computational efforts one should take. Thus, we compute a set of itemsets for a simpler projection, then we remove unpromising itemsets, extend our dataset and the found itemsets with more items (to a more detailed projection). This allows us to reduce the size of the pattern space with a simpler projection and lower computational complexity.

3 Σοφια Algorithm

3.1 Anti-monotonicity w.r.t. a Projection

Our algorithm is based on the projection-antimonotonicity. Many interestingness measures for itemsets, e.g., stability [Kuz07], robustness of closed item-

sets [TMC14], or cosine interest [CWW14], are not (anti-)monotonic w.r.t. inclusion order on itemsets. A measure $\mathcal{M}$ is called *anti-monotonic* if for two itemsets $x \sqsubseteq y$, $\mathcal{M}(x) \geq \mathcal{M}(y)$. For instance, support is an anti-monotonic measure w.r.t. itemset inclusion order and it allows for efficient generation of itemsets with support larger than a threshold [AS94, MTV94, PBTL99]. The projection-antimonotonicity is a generalization of standard anti-monotonicity and allows for efficient processing a larger set of interestingness measures.

Given a projection ψ corresponding to the removal of items Y, *preimages* of an itemset X (we assume $X \cap Y = \emptyset$) for ψ is the set of itemsets $\{Z\}$ such that $\psi(Z) = X$. It can be seen that the set of preimages is given by $\texttt{Preimages}(\texttt{Y}) = \{Z \subseteq \mathcal{I} \mid X \subseteq Z \subseteq X \cup Y\}$. In particular X is also a preimage of itself.

An *anti-monotonic measure $\mathcal{M}$ w.r.t. projection ψ* (or just a *projection-antimonotonic measure*) is a measure which does not increase its value on any premiage of any itemset X for ψ. Since any preimage of X is a superset of X, then any anti-monotonic measure is also a projection-antimonotonic measure.

Example 1. *Let us consider the dataset in Figure 1a. If $\mathcal{M}$ is an interestingness measure w.r.t. a projection ψ and ψ removes item i_5, then $\mathcal{M}(\{i_3\}) \geq \mathcal{M}(\{i_3, i_5\})$. However it is **not** necessary that $\mathcal{M}(\{i_3\}) \geq \mathcal{M}(\{i_3, i_4\})$.*

Thus, given a measure $\mathcal{M}$ anti-monotonic w.r.t. a projection ψ, if y is an itemset such that $\mathcal{M}_\psi(y) < \theta$, then $\mathcal{M}(x) < \theta$ for any preimage x of y for ψ. Hence, if, given an itemset y of $\psi(\mathbb{D})$, one can find all itemsets x of $\mathbb{D}$ such that $\psi(x) = y$, it is possible to find the itemsets in $\psi(\mathbb{D})$ and then to prune them w.r.t. $\mathcal{M}_\psi$, and finally to compute the preimages of the pruned set of itemsets only. It allows one to earlier cut unpromising branches of the search space or adjust a threshold for finding only a limited number of best itemsets.

However, given just one projection, it can be hard to efficiently discover the best itemsets, since the projection is either hard to compute or the number of unpromising itemsets that can be pruned is not high. Corespondingly we need *a chain of projections* $\psi_0 < \psi_1 < \cdots < \psi_k = \mathbb{1}$, where concepts for $\psi_0(\mathbb{D})$ can be easily computed and $\mathbb{1}$ is the identity projection, i.e., $(\forall x)\mathbb{1}(x) = x$. For example, to find frequent itemsets, we typically search for small frequent itemsets and then extend them to larger ones. It corresponds to the extension to a more detailed projection. In particular for binary dataset a chain of projections can be instantiated as a consequent update of a binary dataset with new items.

Chain of projections is a generalization of accessible system [BHPW10]. Given a set of items $\mathcal{I}$ and a subset of its powerset $\mathcal{F} \subseteq 2^{\mathcal{I}}$, the system $(\mathcal{I}, \mathcal{F})$ is accessible if $\forall X \in \mathcal{F} \setminus \{\emptyset\}$ there is $i \in \mathcal{I}$ such that $X \setminus \{i\} \in \mathcal{F}$. Any constraint (or measure) on $2^{\mathcal{I}}$ produces a system of sets. If this system is accessible, then the measure is locally anti-monotonic.

Proposition 1. *A chain of projections can be represented as a sequence of systems $(\mathcal{I}_i, \mathcal{F}_i)$ such that $\mathcal{I}_i \subset \mathcal{I}_{i+1}$ and any element $x \in \mathcal{F}_{i+1}$ is either (1) $x \in \mathcal{F}_i$, or (2) $\exists e \in \mathcal{I}_{i+1} \setminus \mathcal{I}_i$ such that $(x \setminus \{e\}) \in \mathcal{F}_i$, (3) or x accesible in $\mathcal{F}_{i+1}$.*

```
   Data: A dataset 𝔻, a chain of projections Ψ = {ψ_0, ψ_1, ··· , ψ_k}, an anti-monotonic
         measure 𝓜 for the chain Ψ, and a threshold θ for 𝓜.
1  Function ExtendProjection(i, θ, 𝒫_{i−1})
   |  Data: i is the projection number to which we should extend (0 < i ≤ k), θ is a
   |        threshold value for 𝓜, and 𝒫_{i−1} is the set of itemsets for the projection
   |        ψ_{i−1}.
   |  Result: The set 𝒫_i of all itemsets with the value of measure 𝓜 higher than the
   |          threshold θ for ψ_i.
2  |  𝒫_i ⟵ ∅;
3  |  foreach p ∈ 𝒫_{i−1} do
4  |  |  𝒫_i ⟵ 𝒫_i ∪ Preimages(i,p)
5  |  foreach p ∈ 𝒫_i do
6  |  |  if 𝓜_{ψ_i}(p) ≤ θ then
7  |  |  |  𝒫_i ⟵ 𝒫_i \ {p}
8  Function Algorithm_θ-Σοφια
   |  Result: The set 𝒫 of all itemsets with a value of 𝓜 higher than the threshold θ
   |          for 𝔻.
9  |  𝒫 ⟵ FindPatterns(θ, ψ_0);
10 |  foreach 0 < i ≤ k do
11 |  |  𝒫 ⟵ ExtendProjection(i, θ, 𝒫);
```

Algorithm 1: θ-Σοφια

Proof. (1) by idempotency of projections, (2) by contractivity, (3) for deletion of several items. □

3.2 Algorithms

```
   Data: A dataset 𝔻, a chain of projections Ψ = {ψ_0, ψ_1, ··· , ψ_k}, an anti-monotonic
         measure 𝓜 for the chain Ψ, and a threshold L for the maximal number of
         preserved itemsets.
1  Function Algorithm_Σοφια
   |  Result: The threshold θ ensuring that the cardinality of the set 𝒫 is bounded by
   |          L in any step of the algorithm. The set 𝒫 of all itemsets with the value
   |          of measure 𝓜 higher than the threshold θ.
2  |  θ ⟵ θ_min 𝒫 ⟵ FindPatterns(ψ_0);
3  |  foreach 0 < i ≤ k do
4  |  |  θ ⟵ AdjustTheta(θ, L, 𝒫);
5  |  |  𝒫 ⟵ PrunePatterns(θ, 𝒫);
6  |  |  𝒫 ⟵ ExtendProjection(i, θ, 𝒫);
```

Algorithm 2: Σοφια for finding itemsets in $\mathbb{D}$ with the bounded cardinality of the set $\mathcal{P}$.

Given a dataset $\mathbb{D}$ and a measure anti-monotonic w.r.t. a chain of projections, if we are able to find all preimages of any element in the fixed set of a projection ψ_i that belong to a fixed set of the next projection ψ_{i+1}, then we can find all itemsets of the dataset $\mathbb{D}$ with a value of $\mathcal{M}$ higher than a given threshold θ. We call the respective algorithm θ-Σοφια (Algorithm 1). In line 9 we find all itemsets for the dataset $\psi_0(\mathbb{D})$ satisfying the constraint w.r.t. the measure $\mathcal{M}$. Then in lines 10-11 we iteratively extend projections from simpler to more

detailed ones. The extension is done by constructing the set $\mathcal{P}_i$ of preimages of the set $\mathcal{P}_{i-1}$ (lines 2-4) and then by removing the itemsets that do not satisfy the constraint from $\mathcal{P}_i$ (lines 5-7). This listing provides a sketch of the algorithm omitting possible engineering improvements for the sake of simplicity. Most of the known improvements are applicable here, i.e., the ones from [UKA05]. In particular, a canonical order on itemsets is used for mining closed pattern (the theoretical basis for such mining is given in next subsections).

The algorithm is sound and complete, since first, an itemset p is included into the set of preimages of p (since $\psi(p) = p$) and second, if $\mathcal{M}(p) < \theta$, then we remove the itemset p from the set $\mathcal{P}$ and the measure value of any preimage of p is less than θ by the projection chain anti-monotonicity of $\mathcal{M}$. The worst case time complexity for the general case of patterns of θ-Σοφια algorithm is

$$\begin{aligned}\mathbb{T}(\theta\text{-}\Sigma o\varphi\iota\alpha) = \mathbb{T}(FindPatterns(\psi_0))+ \\ + k \cdot \max_{0<i\leq k} |\mathcal{P}_i| \cdot (\mathbb{T}(Preimages) + \mathbb{T}(\mathcal{M})),\end{aligned} \tag{2}$$

where k is the number of projections in the chain, $\mathbb{T}(\mathcal{X})$ is the time for computing the operation $\mathcal{X}$. Since projection ψ_0 can be chosen to be very simple, in a typical case the complexity of $FindPatterns(\theta, \psi_0)$ can be low or even constant. The complexities of $Preimages$ and $\mathcal{M}$ depend on the measure, the chain of projections, and the kind of patterns. In many cases $\max\limits_{0<i\leq k} |\mathcal{P}_i|$ can be exponential in the size of the input, because the number of patterns can be exponential. It can be a difficult task to define the threshold θ such that the maximal cardinality of $\mathcal{P}_i$ is not larger than a given number. Thus, we introduce Σοφια algorithm (Algorithm 2), which automatically adjusts threshold θ ensuring that $\max\limits_{0<i\leq k} |\mathcal{P}_i| < L$. Here L can be considered as a constraint on the memory used by the algorithm. The only difference of Σοφια w.r.t. θ-Σοφια is that after performing an operation that changes the set $\mathcal{P}$ (lines 2 and 6 in Algorithm 2) it adjusts θ in such a way that the cardinality of $\mathcal{P}$ does not exceed the parameter L. It can be seen from (2) that Σοφια has polynomial time complexity if $\mathcal{M}$ and $Preimages$ are polynomial. Indeed, according to (1) if a projection removes only one item, the cardinality of `Preimages` is always 2. Thus, the worst case complexity for θ-Σοφια is

$$\mathbb{T}(\theta\text{-}\Sigma o\varphi\iota\alpha_{\text{binary}}) = |\mathcal{I}| \cdot \max_{0<i\leq N} |\mathcal{P}_i| \cdot \mathbb{T}(\mathcal{M}). \tag{3}$$

We notice that every $\mathcal{P}_i$ is a solution for the projected dataset. Thus, this algorithm has incremental polynomial delay. However, if we fix the available memory L, the complexity of Σοφια for binary data is $|\mathcal{I}| \cdot L \cdot \mathbb{T}(\mathcal{M})$, i.e., it becomes input polynomial modulo complexity of the measure.

To wrap up, in this subsection we have introduced an algorithm for finding top-K itemsets in polynomial time. It is important to notice that the found set of itemsets is exactly the best itemsets w.r.t. to Δ-measure and should not be mixed up with an approximation.

Efficiency Considerations

Recently much work have been done in finding good strategies of enumerating (closed) patterns. Most of them start from the smallest patterns and then iteratively generate larger patterns. It can be naturally expressed as a chain of functions ψ_i that are contractive ($\psi_i(X) \sqsubseteq X$) and idempotent ($\psi_i(\psi_i(X)) = \psi_i(X)$). These functions can be ordered by inclusion of fixed sets because of idempotency. Since these functions are contractive, only patterns larger than a pattern X are preimages of X. Thus, most of the approaches for itemset mining can be formalized by means of a chain of such functions. However, in this work we require a chain of projections, i.e., functions ψ_i, to be also monotonic. It allows us to efficiently mine robust and stable patterns discussed in Section 4. This additional monotonicity still allows one to formalize developed approaches for itemset mining as a chain of projections. However, in this work we does not discuss this formalization and focus on the efficient mining of patterns for nonmonotonic constraints.

3.3 Σοφια Algorithm for Closed Itemsets

Closed frequent itemsets are widely used as a condensed representation of all frequent itemsets since [PBTL99]. Here we show how one can adapt our algorithm for closed patterns. A closed pattern in $\psi_{i-1}(\mathbb{D})$ is not necessarily closed in $\psi_i(\mathbb{D})$. Indeed, if we take the example in Figure 1, the pattern $\{i_1\}$ is closed in $(T, \{i_1, i_2\}, R_2)$ but no more closed in $(T, \{i_1, i_2, i_3\}, R_3)$. However, the extents (closed tidsets) of $\psi(\mathbb{D})$ are extents of $\mathbb{D}$ [GK01]. Thus, we associate the closed patterns with extents, and then work with extents instead of patterns, i.e., a dataset $\mathbb{D} = (T, (2^{\mathcal{I}}, \cap), \delta)$ is transformed into $\mathbb{D}_C = (T, (D_C, \sqcap_C), \delta_C)$, where $D_C = 2^T$. Moreover, for all $a, b \in D_C$ we have $a \sqcap_C b = t(d(a) \sqcap d(b))$, where t and d operators are computed in $\mathbb{D}$ and $\delta_C(t \in T) = \{t\}$. Hence, every pattern p in D_C corresponds to a closed pattern $d(p)$ in $2^{\mathcal{I}}$. A projection ψ of $\mathbb{D}$ induces a projection ψ_C of $\mathbb{D}_C$, given by $\psi_C(A \subseteq T) = t(\psi(d(A)))$ with t and d computed for $\mathbb{D}$.

In the next section we discuss some measures that are anti-monotonic w.r.t. a projection (rather than just anti-monotonic). In the end of the next section we provide an example of how Σοφια works.

4 Itemset Constraints

Table 1: Values of different measures for closed itemesets of context in Figure 1a.

Itemset X	Cosine	$\texttt{Stab}(X)$	$\underset{\alpha=0.9}{\texttt{Rbst}}(X)$	$\Delta(X)$
$\emptyset$	$+\infty$	0.47	0.89991	1
$\{\mathbf{i_3}\}$	**1**	**0.69**	**0.9963**	3
$\{i_1, i_3\}$	0.5	0.5	0.9	1
$\{i_2, i_3\}$	0.5	0.5	0.9	1
$\{i_3, i_4\}$	0.5	0.5	0.9	1
$\{i_3, i_5\}$	0.5	0.5	0.9	1
$\{i_6\}$	1	0.5	0.9	1

4.1 Cosine Interest of an Itemset

The first projection-antimonotonic measure we consider is cosine interest [CWW14]. It is defined by

$$\texttt{Cosine}(X) = \frac{|t(X)|}{\sqrt[|X|]{\prod_{i \in X} |t(\{i\})|}}, \tag{4}$$

i.e., a cosine interest of X is the support of X over the geometric mean of supports of single items from X. As the authors of [CWW14] have shown this measure is not (anti-)monotonic. Then, they also have shown that if we traverse the search space from less supported items to more supported items the cosine interest never decreases. Indeed, given an itemset X and an item i such that $i \notin X$ and $(\forall j \in X)|t(\{i\})| \geq |t(\{j\})|$, we can see that $\texttt{Cosine}(X) \geq \texttt{Cosine}(X \cup \{i\})$ since the itemset support cannot increase while the geometric mean cannot decrease in this case.

To work with cosine interest we can define a projection chain that adds items from less supported ones to more supported, i.e., ψ_1 corresponds to removal of all but the least frequent item from the dataset, ψ_2 corresponds to removal of all but two least frequent items and so on. Then, cosine interesting itemsets can be mined by Σοφια. However this measure is locally anti-monotonic, in the next subsection we consider two proper nonmonotonic measures.

4.2 Stability and Robustness of an Itemset

Stability [Kuz07] and robustness [TMC14] are similar measures when applied to closed itemsets. They measure independence of an itemset w.r.t. subsampling. Stability can only be applied to closed itemsets, while robustness is defined for any type of itemset constraints (closed itemsets, generators, *etc.*). However, in case of closed itemsets neither of them is (anti-)monotonic. Indeed, when robustness is based on an anti-monotonic constraint, it is anti-monotonic. However, closedness of itemsets is not an anti-monotonic constraint. Since stability and robustness are similar, we define them on a similar basis.

Given a dataset $\mathbb{D} = (T, \mathcal{I}, R)$, a triple $(S, \mathcal{I}, R)$ where $S \subseteq T$ is called a *subdataset* of $\mathbb{D}$. If we give a weight to every subdataset of $\mathbb{D}$, then we can find the sum of weights of all subdatasets of $\mathbb{D}$ where an itemset X is closed. This sum gives us stability or robustness of the closed itemset X depending on how we define the weights of subdatasets.

In the case of *stability* the weights w of all subdatasets $\mathbb{D}_s$ of $\mathbb{D}$ are equal, i.e., $w(\mathbb{D}_s) = 2^{-|T|}$. In this case we consider every subdataset equally probable and compute the probability that the itemset X is closed.

Example 2. *Consider example in Figure 1a. The set of concepts (the pattern of every concept is a closed itemset) is shown in Figure 1b. Stability of every closed itemset is shown in Table 1. Let us consider the highlighted itemset* $X = \{i_3\}$. *There are* 2^5 *possible subdatasets. Only in the following 10 subdatasets* X *is not closed (only the set of transactions for every subdataset is given):*

$\emptyset, \{t_1\}, \ldots, \{t_5\}, \{t_1, t_5\}, \{t_2, t_5\}, \{t_3, t_5\}, \{t_4, t_5\}$. *Thus, stability of* X *can be found as* $\texttt{Stab}(X) = 1 - 10 \cdot 2^{-5} = 0.69$.

It should be noticed that stability of all comparable itemsets in the lattice is smaller than stability of X*, which highlights the nonmonotonicity of stability.*

In the case of *robustness* the weights w of subdatasets are computed differently. These weights depend on a parameter $0 \leq \alpha \leq 1$ denoting the probability of a transaction to be retained in the dataset. The weight of a subdataset $\mathbb{D}_s = (S, \mathcal{I}, R)$ of $\mathbb{D} = (T, \mathcal{I}, R)$ corresponds to the probability of obtaining $\mathbb{D}_s$ by removing every single transaction from $\mathbb{D}$ with probability $1 - \alpha$: $w(\mathbb{D}_s) = \alpha^{|S|} \cdot (1 - \alpha)^{|T|-|S|}$.

Example 3. *Consider example in Figure 1a. Robustness for* $\alpha = 0.9$ *for every closed itemset is shown in Table 1. Let us consider the highlighted itemset* $X = \{i_3\}$*. It is not closed in the same as above 10 subdatasets but their weights are different (the weights are shown in superscripts):* $\emptyset^{w=10^{-5}}, \{t_1\}^{w=9\cdot 10^{-5}}, \ldots, \{t_5\}^{w=9\cdot 10^{-5}}, \{t_1, t_5\}^{w=8.1\cdot 10^{-4}}, \{t_2, t_5\}^{w=8.1\cdot 10^{-4}}, \{t_3, t_5\}^{w=8.1\cdot 10^{-4}}, \{t_4, t_5\}^{w=8.1\cdot 10^{-4}}$*. Thus, robustness of* X *for* $\alpha = 0.9$ *is equal to* $\texttt{Rbst}_{\alpha=0.9}(X) = 0.9963$*. It can be verified that robustness is not an anti-monotonic measure.*

It is not hard to show that independently of the weights w of subdatasets, stability and robustness are anti-monotonic measures w.r.t. any projection.

Proposition 2. *Stability and robustness are anti-monotonic measures w.r.t. any projection.*

Proof. Here we want to show that for any projection ψ if a pattern X is closed in a subdataset $\mathbb{D}_s$ then $\psi(X)$ is closed in $\psi(\mathbb{D}_s)$, where $\mathbb{D}_s = (S, \mathcal{I}, R)$ is a subdataset of $\mathbb{D} = (T, \mathcal{I}, R)$ with $S \subseteq T$. We note that if X is closed in $\mathbb{D}_s$ it is also closed in $\mathbb{D}$. And since $\psi(X)$ closed in $\psi(\mathbb{D})$, then for projection ψ_C from Section 3.3 we have $d(\psi_C(t(X))) = \psi(X)$. Hence, we can work with images of ψ on closed patterns in order to find the corresponding images of ψ_C.

Let $Y = d(t(\psi(X)) \cap S)$ be a closure of $\psi(X)$ in $\mathbb{D}_s$. Since $\psi(X) \sqsubseteq X$, then $t(\psi(X)) \supseteq t(X)$. Hence $S \cap t(\psi(X)) \supseteq S \cap t(X)$. Then $Y = d(S \cap t(\psi(X))) \sqsubseteq d(S \cap t(X)) = X$, since Y is the closure of $\psi(X)$ in $\psi(\mathbb{D}_s)$ and X is the closure of X in $\mathbb{D}_s$. Thus, we have $Y \sqsubseteq X$ and $c_t(Y) \sqsubseteq c_t(X) = X$. Because of monotonicity of projections one has $\psi(c_t(Y)) \sqsubseteq \psi(X)$ and hence $Y \sqsubseteq \psi(X)$.

Since Y is the closure of $\psi(X)$ in $\mathbb{D}_s$, then $Y \sqsupseteq \psi(X)$. Hence $Y = \psi(X)$. □

Estimates of Stability and Robustness

For both stability and robustness it is shown that the corresponding constraint is NP-hard [Kuz07, TMC14]. Thus, for efficient mining, estimates of stability and robustness are essential. Here we introduce a fast computable estimate of robustness in the same way we did it for stability in [BKN14].

Let us consider closed itemsets X and Y such that $X \subset Y$. *Can we define the subdatasets where* X *is not closed?* Let us define $\Delta(X, Y)$ as the cardinality of the set of transactions described by X but not by Y: $\Delta(X, Y) = t(X) \setminus t(Y)$.

This set is not empty since $X \neq Y$ and they are closed. It is clear that X is not closed in any subdataset that removes all transactions from $\Delta(X,Y)$, since Y is a larger itemset with the same support. Then, $\texttt{Stab}(X) \leq 1 - 2^{-\Delta(X,Y)}$ and $\texttt{Rbst}(X) \leq 1 - (1-\alpha)^{\Delta(X,Y)}$ for any closed itemset $Y \supset X$. In particular, we can put Y to the closest closed superitemset of X.

In the same way we can take all immediate closed superitemsets of X and take into account all the subdatasets where X is not closed. Since some of the subdatasets are probably counted several times we get the lower bound, i.e., $\texttt{Stab}(X) \geq 1 - \sum\limits_{Y \prec X} 2^{-\Delta(X,Y)}$ and $\texttt{Rbst}(X) \geq 1 - \sum\limits_{Y \prec X} (1-\alpha)^{\Delta(X,Y)}$.

Proposition 3. *Stability and robustness are bounded as follows, where $X \prec Y$ means that X is an immediate closed subitemset of Y:*

$$1 - \sum_{Y \prec X} 2^{-\Delta(X,Y)} \leq \texttt{Stab}(X) \leq 1 - 2^{-\Delta(X,Y)} \tag{5}$$

$$1 - \sum_{Y \prec X} (1-\alpha)^{\Delta(X,Y)} \leq \texttt{Rbst}(X) \leq 1 - (1-\alpha)^{\Delta(X,Y)} \tag{6}$$

In particular we can see that when $\alpha = 0.5$ the estimates are exactly the same. As it is recently shown [BKN14], the estimate of stability is quite precise for the concepts with stability close to 1. Then, when $\alpha > 0.5$ the precision of the estimate of robustness is even more precise.

These estimates can be computed in polynomial time in contrast to stability and robustness. And thus we can use one of the bounds as a proxy to stability and robustness. It can be seen that the rankings based on the upper bound of stability and robustness are exactly the same as the ranking based on $\Delta(X) = \min\limits_{Y \prec X} \Delta(X,Y)$. Although for the lower bound of stability and robustness it is hard to show the projection anti-monotonicity, we can show it for the upper bound. In the following $\Delta(X)$ is called Δ-measure.

Proposition 4. *Δ-measure is an anti-monotonic measure w.r.t. any projection.*

Proof. We remind that for dealing with closed patterns the tidsets are considered as patterns as discussed in Section 3.3. By properties of projections, if an extent (the tidset of a concept) is found in $\psi(\mathbb{D})$, it is necessarily found in $\mathbb{D}$ [GK01]. Let us consider a tidset E of a concept and a tidset of its descendant E_c in $\psi(\mathbb{D})$, where a descendant concept is a concept with a smaller tidset and larger itemset. Let us suppose that E_p is a preimage of E for the projection ψ. Since E_c and E_p are extents in $\mathbb{D}$, the set $E_{cp} = E_c \cap E_p$ is an extent in $\mathbb{D}$ (the intersection of two closed sets is a closed set). Since E_p is a preimage of E, then $E_p \not\subseteq E_c$ (otherwise, E_p is a preimage of E_c and not of E). Then, $E_{cp} \neq E_p$ and $E_{cp} \subseteq E_p$. Hence, $\Delta(E_p) \leq |E_p \setminus E_{cp}| \leq |E \setminus E_c|$. So, given a preimage E_p of E, $(\forall E_c < E)\Delta(E_p) \leq |E \setminus E_c|$, i.e., $\Delta(E_p) \leq \Delta(E)$. Thus, we can use Δ-measure in combination with Σοφια algorithm. □

Example 4. *Consider example in Figure 1a. Δ-measure for every closed itemset is shown in Table 1. Let us consider the highlighted itemset $X = \{i_3\}$ with support equal to 4. The closest superitemsets of X are $\{i_1, i_3\}$, $\{i_2, i_3\}$, $\{i_3, i_4\}$,*

Table 2: Patterns given by their extent and their stability in the contexts corresponding to a chain of projections.

#	Pattern Ext.	Δ-measure						
		$\mathcal{I}_0$	$\mathcal{I}_1$	$\mathcal{I}_2$	$\mathcal{I}_3$	$\mathcal{I}_4$	$\mathcal{I}_5$	$\mathcal{I}_6$
1	12345	5	4	4	1	1	1	1
2	1	–	1	1	1	1	1	1
3	2	–	–	1	1	1	1	1
4	1234	–	–	–	3	3	3	3
6	3	–	–	–	–	1	1	1
7	4	–	–	–	–	–	1	1
8	5	–	–	–	–	–	–	1

and $\{i_3, i_5\}$*, all having support equal to one. Thus,* Δ*-measure of* X *is equal to* $\Delta(X) = 4 - 1 = 3$*. It can be noticed that* Δ*-measure is not an (anti-)monotonic measure.*

Δ-measure is related to the work of margin-closeness of an itemset [MTU11]. In this work, given a set of patterns, e.g., frequent closed patterns, the authors rank them by the minimal distance in their support to the closest superpattern divided by the support of the pattern. In our case, the minimal distance is exactly the Δ-measure of the pattern.

4.3 Example of Stable Itemsets in Binary Data

Let us consider the example in Figure 1 and show how we can find all Δ-stable itemsets with threshold $\theta = 2$. We have a binary dataset $\mathbb{D} = (T, \{i_1, \cdots, i_6\}, R)$. Let us denote $\mathcal{I}_i = \{i_1, \cdots, i_i\}$. The sets $\mathcal{I}_i$ correspond to a chain of projections.

In Table 2 all closed itemsets are given by the corresponding tidsets, i.e., by elements of D_C. For simplicity we write 1234 instead of $\{t_1, t_2, t_3, t_4\}$. For every element Δ-measure is shown for every $\mathcal{I}_i$. A cell is shown in gray if the itemset is no more considered (the value of Δ is less than 2).

For example, in the transition from $\mathcal{I}_2$ to $\mathcal{I}_3$ the set 1234 is discovered with $\Delta(1234) = 3$, but $\Delta(12345) = 5 - 4 = 1$ which is less than $\theta = 2$. Thus, itemset 12345 is discarded and highlighted gray. The global process is as follows (for the example in Figure 1). In the empty binary dataset $(T, \emptyset, R)$ the first itemset 12345 is considered. Then, in $(T, \{i_1\}, R)$ a possible preimage of 12345 can be either 12345 or $12345 \cap t(\{i_1\}) = 1$. The set 12345 is Δ-stable ($\Delta(12345) = 4$), while 1 is not Δ-stable ($\Delta(1) = 1$) and is discarded. Then, the process continues with $(T, \{i_1, i_2\}, R)$ and 12345 is kept while $12345 \cap t(\{i_2\}) = 2$ is removed for the same reason as 1. After that, with $(T, \{i_1, i_2, i_3\}, R)$ two preimages are still considered, 12345 and 1234. This time $\Delta(1234) = 3$, while $\Delta(12345) = 1$

Table 3: Computational efficiency of Σοφια algorithm.

Dataset	Top-K	$\theta_{Supp.}$	θ_{Δ}	LCMv3	Δ	Charm-L ($\sim$ Charm-L + Δ)	Σοφια
FIMI							
chess	3	1145	234	1.62	> 100	> 100	**0.03**
	928	277	98	> 100	—	> 100	**0.13**
connect	1	25466	4224	0.21	128	111	**0.61**
	1000	8822	2602	1.25	> 100	> 100	**1.77**
mushroom	1	6272	2256	< 0.01	0.07	**0.01**	0.05
	722	216	193	0.06	2.12	0.50	**0.23**
pumsb	1	33128	2035	0.15	> 300	36.7	**0.8**
	984	8793	865	> 300	—	> 300	**38.7**
pumsb*	1	30787	8090	0.04	1.42	0.16	**0.65**
	997	2808	834	4.47	> 300	> 300	**27.8**
LUCS							
adult	1	34338	6939	0.01	0.78	0.05	**0.20**
	998	674	446	0.11	16.45	2.15	**1.27**
waveform	1	3424	1179	$<$ **0.01**	**0.01**	**0.01**	0.03
	984	401	141	0.09	4.42	1.24	**0.25**
UCI							
plants	1	11676	6154	< 0.01	0.11	0.02	**0.11**
	984	649	148	> 100	—	> 100	**0.96**

and the set 12345 is discarded. The process continues in the same way with $\Delta(1234) = 3$ and all other possible elements are discarded.

5 Experiments and Discussion

5.1 Comparing Computational Efficiency

In the first experiment we show the computational efficiency of Σοφια coded in C++[1]. We use public available big datasets from FIMI[2], LUCS [Coe03], and UCI [FA10] repositories. The experiments are carried out on an "Intel(R) Core(TM) i7-2600 CPU @ 3.40GHz" computer with 8Gb of memory under Ubuntu 14.04.

We should note two points here. First, to the best of our knowledge Σοφια is the first algorithm that computes top Δ-stable and robust itemsets, so there are no direct competitors. Moreover, computing Δ-measure for an itemset requires either a known partial order of itemsets or a search for its descendants (closed supersets). Thus, as an approximate competitors we decided to use two algorithms `LCMv3` [UKA05] and `Charm-L` [ZH05]. The first one is one of the most efficient algorithm for itemset mining that should be followed by Δ-measure computation for every concept. `Charm-L` is less efficient than `LCMv3`, but allows

[1] The implementation is available at `https://github.com/AlekseyBuzmakov/FCAPS`

[2] `http://fimi.ua.ac.be/data/`

one to find the partial order of itemsets necessary for the fast computation of Δ-measure.

Second, the current implementation of $\Sigma o\varphi\iota\alpha$ does not use most of the modern optimization techniques, e.g., like in `LCMv3` [UKA05]. The current implementation relies only on the so-called conditional database, i.e., where for every tidset X the items that belong to all transactions from $t(X)$ and the items that belong to neither transactions from $t(X)$ are recorded [UKA05]. But nevertheless, the computation with the current implementation is efficient.

The experiment is organized as following. First, $\Sigma o\varphi\iota\alpha$ finds around the 1000 most Δ-stable itemsets and the maximal support threshold ensuring to find all these the most Δ-stable itemsets. Among them we find the most Δ-stable itemset (or itemsets if they have the same value of Δ-measure) and the corresponding support threshold. So `LCMv3` and `Charm-L` are additionally provided with an oracle returning the required support thresholds. For these two thresholds we run `LCMv3` and `Charm-L` algorithm and register the computation time. In addition for `LCMv3` we register also the time needed for computing Δ-measure, while for `Charm-L` this time is insignificant. In Table 3 for every dataset we give the results corresponding to every threshold, and the corresponding thresholds for support and Δ-measure. For example, for dataset `chess` we run two experiments. In the first one we search for top-3 Δ-stable itemsets having the same value (234) for Δ-measure. The less frequent itemsets among these three has support equal to 1145, thus, `LCMv3` and `Charm-L` should be run with this support threshold in order to enumerate all of these itemsets. `LCMv3` finds the corresponding frequent closed itemset in 1.67 seconds, then it takes more than 100 seconds for computing Δ-measure. `Charm-L` takes more than 100 seconds and $\Sigma o\varphi\iota\alpha$ requires only 0.03 seconds. In the second experiment for dataset `chess` we search for top-928 Δ-stable itemsets, all of them have support at least 277 and Δ-measure 98.

We boldify the computation time for an algorithm in Table 3, if it is better than the time of the competitors. We can see that even `LCMv3` alone does not always beat $\Sigma o\varphi\iota\alpha$, while the additional time for `LCMv3` for computing Δ-measure is always significant. There are only two cases when $\Sigma o\varphi\iota\alpha$ is slightly worse (FIMI-mushroom and LUCS-waveform). For both cases the most stable itemset has a very high support and only a couple of itemsets are frequent enough in both datasets. In contrast, if the frequency of the most Δ-stable itemsets is not high, then $\Sigma o\varphi\iota\alpha$ is many times faster than even `LCMv3` alone.

In these experiments we do not provide the found itemsets since the main focus of our paper is efficiency. However, we highlight that Δ-stable patterns are not trivial and can be found deep in the lattice of patterns [MLB$^+$15].

5.2 Scalability

We can study scalability of $\Sigma o\varphi\iota\alpha$ from different points of view. First, we can measure the time necessary for finding top-L concepts, i.e., how the memory limitation L changes the efficiency. It is shown in Table 4 for the same datasets. We can see that the computation time changes linearly w.r.t. the memory limitation L as it is expected from Eq. (3).

Table 4: Scalability of Σοφια w.r.t. the number of stored patterns

Dataset	L=100	L=1000	$L = 10^4$	$L = 10^5$
		FIMI		
chess	0.04	0.13	1.35	14.7
connect	0.70	1.77	12.7	131
mushroom	0.1	0.29	2.62	40.5
pumsb	7.15	71.5	904	—
pumsb*	4.14	45.7	832	—
		LUCS		
adult	0.30	0.99	8.79	83.97
waveform	0.06	0.18	1.97	22.13
		UCI		
plants	0.22	1.09	11.58	117.91

Finally, we check how computation time depends on the size of the dataset. For that we run our experiments for $L = 1000$, and vary the number of transactions in a dataset. We permute several time the order of transactions of the dataset. For every permutation we construct datasets containing certain amount (the size of the dataset) of the first transactions from this permutation. The computation time is averaged over the permutations. Figure 2 shows the computation time necessary to process a certain fraction of transactions in the dataset. Time is given as a fraction of time for processing the whole dataset. We can see that computation time changes linearly w.r.t. the fraction of processed transactions.

Σοφια allows limiting the memory in use; thus, as long as the program fits within the memory, which is controllable, the scalability of our approach is linear w.r.t. to the size of the dataset, and consequently can be applied to very huge datasets.

6 Conclusion

In this paper we have introduced a new class of interestingness measures, so-called projection-antimonotonic measures. This wide class of measures includes classical anti-monotonic, locally anti-monotonic, and some nonmonotonic measures. We have introduced algorithm Σοφια, which allows one to efficiently mine patterns w.r.t. projection-antimonotonic measures. We have studied stability and robustness, two projection-antimonotonic measures, and have introduced polynomial estimates of them, called Δ-measure. Finally, in the experimental part we have showed that Σοφια can find Δ-stable itemsets much more efficiently than postpruning approaches.

Many directions for future work are promising. First, we should work on adaptation of Σοφια for dealing with different kinds of pattern structures, e.g., based on sequences or graphs [KS05]. Second, Σοφια allows one to introduce new data mining approaches by means of projections of special kind, thus, it is interesting to study possible classes of projections. Finally, besides robustness

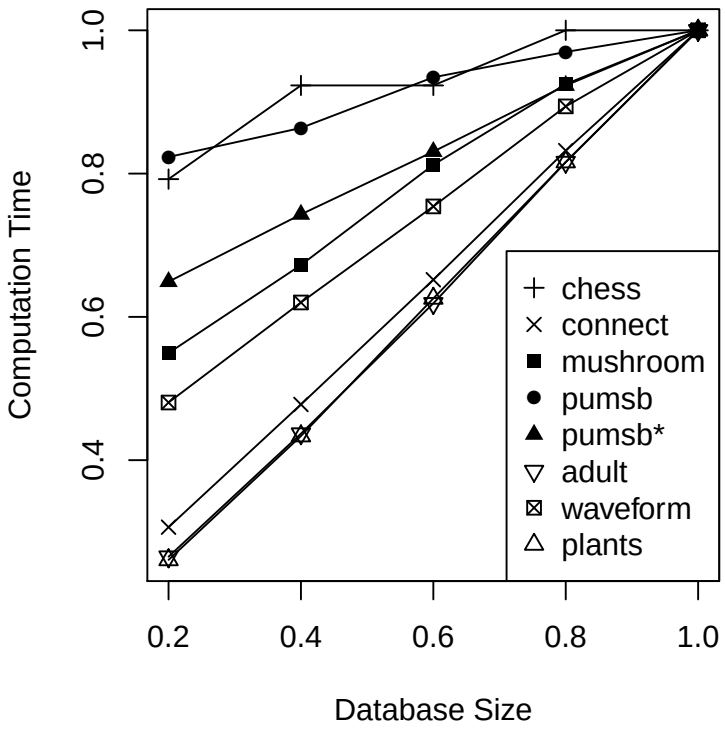

Figure 2: Scalability of Σοφια w.r.t. dataset size. X-axis shows the fraction of objects taken from an original dataset, and Y-axis shows the fraction of time w.r.t. the computational time needed for processing the original dataset.

and stability we a study of other projection-antimonotonic measures is important.

References

[AS94] Rakesh Agrawal and Ramakrishnan Srikant. Fast algorithms for mining association rules. In *Proc. 20th int. conf. very large data bases, VLDB*, volume 1215, pages 487–499, 1994.

[BHPW10] Mario Boley, Tamás Horváth, Axel Poigné, and Stefan Wrobel. Listing closed sets of strongly accessible set systems with applications to data mining. *Theor. Comput. Sci.*, 411(3):691–700, jan 2010.

[BKN14] Aleksey Buzmakov, Sergei O. Kuznetsov, and Amedeo Napoli. Scalable Estimates of Concept Stability. In Christian Sacarea, Cynthia Vera Glodeanu, and Mehdi Kaytoue, editors, *Form. Concept Anal.*, volume 8478 of *LNCS*, pages 161–176. Springer Berlin Heidelberg, 2014.

[BKN15] Aleksey Buzmakov, Sergei O. Kuznetsov, and Amedeo Napoli. Fast Generation of Best Interval Patterns for Nonmonotonic Constraints. In Annalisa Appice, Pedro Pereira Rodrigues, Vítor Santos Costa, João Gama, Alípio Jorge, and Carlos Soares, editors, *Mach. Learn. Knowl. Discov. Databases*, volume 9285 of *LNCS*, pages 157–172. Springer International Publishing, 2015.

[Coe03] F. Coenen. *The LUCS-KDD Discretised and normalised ARM and CARM Data Library*[3]. Department of Computer Science, The University of Liverpool, UK, 2003.

[CRB08] Loïc Cerf, Céline Robardet, and Jean-François Boulicaut. Data-Peeler: Constraint-Based Closed Pattern Mining in n-ary Relations. In *SDM'08 Proc. Eighth SIAM Int. Conf. Data Min.*, pages 37—-48. SIAM, 2008.

[CWW14] Jie Cao, Zhiang Wu, and Junjie Wu. Scaling up cosine interesting pattern discovery: A depth-first method. *Inf. Sci. (Ny).*, 266(0):31–46, 2014.

[FA10] A. Frank and A. Asuncion. *UCI Machine Learning Repository [http://archive.ics.uci.edu/ml].* University of California, Irvine, School of Information and Computer Sciences, 2010.

[GK01] Bernhard Ganter and Sergei O. Kuznetsov. Pattern Structures and Their Projections. In Harry S. Delugach and Gerd Stumme, editors, *Concept. Struct. Broadening Base*, volume 2120 of *LNCS*, pages 129–142. Springer Berlin Heidelberg, 2001.

[GW99] Bernhard Ganter and Rudolf Wille. *Formal Concept Analysis: Mathematical Foundations.* Springer, 1st edition, 1999.

[KKN11] Mehdi Kaytoue, Sergei O. Kuznetsov, and Amedeo Napoli. Revisiting Numerical Pattern Mining with Formal Concept Analysis. In *IJCAI 2011, Proc. 22nd Int. Jt. Conf. Artif. Intell. Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 1342–1347, 2011.

[KS05] Sergei O. Kuznetsov and Mikhail V. Samokhin. Learning Closed Sets of Labeled Graphs for Chemical Applications. In Stefan Kramer and Bernhard Pfahringer, editors, *Inductive Log. Program. SE - 12*, volume 3625 of *LNCS*, pages 190–208. Springer Berlin Heidelberg, lecture no edition, 2005.

[Kuz07] Sergei O. Kuznetsov. On stability of a formal concept. *Ann. Math. Artif. Intell.*, 49(1-4):101–115, 2007.

[3] `http://www.csc.liv.ac.uk/~frans/KDD/Software/LUCS_KDD_DN/`

[MLB+15] Jean-Philippe Métivier, Alban Lepailleur, Aleksey Buzmakov, Guillaume Poezevara, Bruno Crémilleux, Sergei Kuznetsov, Jérémie Le Goff, Amédéo Napoli, Ronan Bureau, and Bertrand Cuissart. Discovering structural alerts for mutagenicity using stable emerging molecular patterns. *J. Chem. Inf. Model.*, 55(5):925–940, apr 2015.

[MTU11] Fabian Moerchen, Michael Thies, and Alfred Ultsch. Efficient mining of all margin-closed itemsets with applications in temporal knowledge discovery and classification by compression. *Knowl. Inf. Syst.*, 29(1):55–80, 2011.

[MTV94] Heikki Mannila, Hannu Toivonen, and A Inkeri Verkamo. Efficient Algorithms for Discovering Association Rules. In *Knowl. Discov. Data Min.*, pages 181–192, 1994.

[PBTL99] Nicolas Pasquier, Yves Bastide, Rafik Taouil, and Lotfi Lakhal. Efficient Mining of Association Rules Using Closed Itemset Lattices. *Inf. Syst.*, 24(1):25–46, 1999.

[ROK08] Camille Roth, Sergei A. Obiedkov, and Derrick G. Kourie. On succinct representation of knowledge community taxonomies with formal concept analysis. *Int. J. Found. Comput. Sci.*, 19(02):383–404, apr 2008.

[SC05] Arnaud Soulet and Bruno Crémilleux. Optimizing constraint-based mining by automatically relaxing constraints. In *Proc. 5th IEEE Inter- Natl. Conf. Data Min. (ICDM 2005)*, pages 777–780. IEEE Computer Society, 2005.

[SC08] Arnaud Soulet and Bruno Crémilleux. Adequate condensed representations of patterns. *Data Min. Knowl. Discov.*, 17(1):94–110, 2008.

[SDB13] Eirini Spyropoulou, Tijl De Bie, and Mario Boley. Interesting pattern mining in multi-relational data. *Data Min. Knowl. Discov.*, (April):1–42, 2013.

[TMC14] Nikolaj Tatti, Fabian Moerchen, and Toon Calders. Finding Robust Itemsets under Subsampling. *ACM Trans. Database Syst.*, 39(3):1–27, 2014.

[UKA05] Takeaki Uno, Masashi Kiyomi, and Hiroki Arimura. LCM Ver.3: Collaboration of Array, Bitmap and Prefix Tree for Frequent Itemset Mining. In *Proc. 1st Int. Work. Open Source Data Min. Freq. Pattern Min. Implementations*, OSDM '05, pages 77–86, New York, NY, USA, 2005. ACM.

[VT14] Jilles Vreeken and Nikolaj Tatti. Interesting Patterns. In Charu C Aggarwal and Jiawei Han, editors, *Freq. Pattern Min.*, pages 105–134. Springer International Publishing, 2014.

[YCHY08] Xifeng Yan, Hong Cheng, Jiawei Han, and Philip S. Yu. Mining significant graph patterns by leap search. In *Proc. 2008 ACM SIGMOD Int. Conf. Manag. data - SIGMOD '08*, pages 433–444, New York, New York, USA, jun 2008. ACM Press.

[ZH05] Mohammed J. Zaki and Ching-Jui Hsiao. Efficient algorithms for mining closed itemsets and their lattice structure. *IEEE Trans. Knowl. Data Eng.*, 17(4):462–478, apr 2005.